

CORRECTED VERSION

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
8 June 2000 (08.06.2000)

PCT

(10) International Publication Number  
**WO 00/033195 A1**

- (51) International Patent Classification<sup>7</sup>: G06F 12/08, 9/46 (74) Agents: TERRILE, Stephen, A. et al.; Skjerven, Morrill, MacPherson, Franklin & Friel, LLP, Suite 700, 25 Metro Drive, San Jose, CA 95110 (US).
- (21) International Application Number: PCT/US99/28875 (81) Designated States (*national*): JP, KR.
- (22) International Filing Date: 3 December 1999 (03.12.1999) (84) Designated States (*regional*): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 09/204,794 3 December 1998 (03.12.1998) US Published:  
— with international search report
- (48) Date of publication of this corrected version: 29 August 2002
- (71) Applicant: SUN MICROSYSTEMS, INC. [US/US]; 901 San Antonio Road, M/S PAL01-521, Palo Alto, CA 94303 (US). (15) Information about Correction:  
see PCT Gazette No. 35/2002 of 29 August 2002, Section II
- (72) Inventors: JOY, William, N.; P.O. Box 23, Aspen, CO 81612-0023 (US). O'CONNOR, James, M.; 1877 Slate Drive, Union City, CA 94587 (US). TREMBLAY, Marc; 140 Hanna Way, Menlo Park, CA 94024 (US). For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.



WO 00/033195 A1

(54) Title: ELIMINATION OF TRAPS AND ATOMICITY IN THREAD SYNCHRONIZATION

(57) Abstract: Elimination of traps and atomics in thread synchronization is provided. In one embodiment, a processor includes a lock cache. The lock cache holds a value that corresponds to or identifies a computer resource only if a current thread executing on the processor owns the computer resource. A lock cache operation (e.g., a lockcachecheck instruction) determines whether a value identifying a computer resource is cached in the lock cache and returns a first predetermined value if the value identifying the computer resource is cached in the lock cache. Otherwise, a second predetermined value is returned.

- 1 -

**ELIMINATION OF TRAPS AND ATOMICS IN THREAD SYNCHRONIZATION****TECHNICAL FIELD**

The present invention relates to microprocessors, and more particularly, to locking of computer resources.

**5      BACKGROUND ART**

When different computer entities such as computer processes or threads share a computer resource (for example, data, code, or a piece of hardware), it may be desirable to allow one of the computer entities to lock a resource for a while to prevent some types of access to the resource by other computer entities. For example, if two or more threads share computer data, and one thread has started but not finished to modify the data when  
10 another thread is accessing the data, the other thread may get incorrect information from the data and/or the data could be corrupted by the two threads. Also, if one thread has started but not finished execution of a critical code section when another thread starts executing the same code section, execution errors may occur if, for example, the critical code section modifies the state of a data area, a hardware controller, or some other computer resource. Therefore, locking techniques have been provided to allow computer entities to lock computer resources.

15            It is desirable to provide fast techniques for locking of computer resources.

**DISCLOSURE OF INVENTION**

The present invention provides elimination of traps and atomics in thread synchronization. Efficient and fast thread synchronization is provided. For example, thread synchronization can be efficient and fast in some frequently occurring situations, such as when a thread owns a computer resource (e.g., has a lock on a computer  
20 resource).

In one embodiment, a lock cache (e.g., four registers) that holds up to four values, each of which is a reference to a locked object (e.g., an address of the locked object) is provided. The lock cache maintains LOCKCOUNTs. A lockcachecheck instruction (e.g., lockcachecheck object\_address, Rd) returns a first predetermined value if the object\_address is stored in the lock\_cache (e.g., the current thread executing on the  
25 processor owns the object). Otherwise, the lockcachecheck instruction returns a second predetermined value (e.g., the lock cache is not currently storing the object\_address, and the current thread may or may not own the object, which can be determined based on an access to memory that is handled in software). For example, if the first predetermined value is returned, then processing can continue with the current instruction stream (the current thread owns the object). If the second predetermined value is returned, then processing can branch to a

- 2 -

different set of instructions (e.g., to check a table in memory to determine whether the current thread owns the object, and if not, then the current thread can wait for the object to become available).

Other features and advantages of the present invention are described below.

#### **BRIEF DESCRIPTION OF DRAWINGS**

5           FIG. 1 is a block diagram of a computer system including a processor according to the present invention.

FIG. 2 is a block diagram showing registers that are used for locking operations in the processor of FIG. 1, and also showing related data structures in the memory of the system of FIG. 1.

FIG. 3 is a block diagram showing data structures in the memory of FIG. 1.

10           FIG. 4 is a block diagram showing registers used for locking operations in a processor according to the present invention.

FIG. 5 is a block diagram illustrating lock caches in accordance with one embodiment of the present invention.

15           FIG. 6 is a functional diagram illustrating the operation of a lockcachecheck instruction in accordance with one embodiment of the present invention.

FIG. 7 is a flow diagram illustrating the operation of a lockcachecheck instruction in accordance with one embodiment of the present invention.

#### **MODES FOR CARRYING OUT THE INVENTION**

20           FIG. 1 is a block diagram of a computer system including locking circuitry. Processor 110 is connected to memory 120 by bus 130. Processor 110 includes execution unit 136 which executes instructions read from memory 120. Execution unit 136 includes lock registers 144 labeled LOCKADDR, LOCKCOUNT. These registers are used for object locking as described below.

25           Bus 130 is connected to I/O bus and memory interface unit 150 of processor 110. When processor 110 reads instructions from memory 120, interface unit 150 writes the instructions to read instruction cache 156. Then the instructions are decoded by decode unit 160. Decode unit 160 sends control signals to execution control and microcode unit 166. Unit 166 exchanges control signals with execution unit 136. Decode unit 160 also sends control signals to stack cache and stack cache control unit 170 (called "stack cache" below) or the register file in processors that do not include a stack cache. Stack cache 170 exchanges control and data signals

- 3 -

with execution unit 136 and data cache unit 180. Cache units 170 and 180 exchange data with memory 120 through interface 150 and bus 130. Execution unit 136 can flush instruction cache 156, stack cache 170 and data cache 180.

FIG. 2 illustrates registers 144 and one of the corresponding objects in memory 120. Registers 144  
5 include four register pairs labeled LOCKADDR0/LOCKCOUNT0 through LOCKADDR3/LOCKCOUNT3. Each LOCKADDR register is to hold an address of a locked object. In a preferred embodiment, each address is 32 bits wide, and accordingly each LOCKADDR register is 32 bits wide. However, in some embodiments, each object starts on a 4-byte boundary. Therefore, in some embodiments the two least significant bits of the object's address are zero, and are omitted from registers LOCKADDR. In such an embodiment, each register  
10 LOCKADDR is 30-bits wide.

If a LOCKADDR register contains 0, this means the register pair is unused.

In each register pair, the LOCKCOUNT register holds the count of lock instructions for the object whose address is held in the corresponding LOCKADDR register. The LOCKCOUNT register holds the number of those lock instructions for which a corresponding unlock instruction has not issued. The LOCKCOUNT  
15 register is incremented on each lock instruction for the object, and is decremented on each unlock instruction. The lock is actually freed only when the LOCKCOUNT register is decremented to zero. (However, in some embodiments, the LOCKCOUNT register holds only a portion of the lock count, as described below. The lock is freed when the entire lock count is decremented to zero.) In some embodiments, each LOCKCOUNT register is 8-bits wide, to hold a number between 0 and 255.

20 Multiple lock instructions without intervening unlock instructions may be a result of recursive code. Because the LOCKCOUNT registers keep the net count of the lock and unlock instructions for the objects (that is, the difference between the numbers of the lock instructions and the unlock instructions for the object), software programs are relieved from the need to do a test before each unlock instruction to determine whether the object was locked by some other part of the thread and should therefore remain locked until the need for that lock  
25 has expired.

In some embodiments, registers 144 keep lock addresses and counts for one thread or one computer process only. When processor 110 switches to a different thread or process, registers 144 are loaded with lock data (lock addresses and counts) for the new thread or process which is to be executed. Accordingly, lock registers 144 allow processor 110 to avoid atomic operations when an address hit occurs. Atomic operations  
30 such as a swap operation, a compare and swap operation and a test and set operation can significantly reduce the performance of the processor.

- 4 -

FIG. 2 illustrates an object whose address is stored in a LOCKADDR register (register LOCKADDR3 in FIG. 2). In FIG. 2, the object is shown stored in memory 120. However, all or part of the object can be stored in data cache 180. Throughout this description, when we describe storing data or instructions in memory 120, it is to be understood that the data or instructions can be stored in data cache 180, stack cache 170 or instruction cache 156, unless mentioned otherwise.

As shown in FIG. 2, the address in register LOCKADDR3 is a pointer to object structure 220. Object structure 220 starts with a header 220H. Header 220H is followed by other data (not shown). Header 220H includes a pointer to class structure 230 describing the object. Class structure 230 is aligned on a 4-byte boundary. As a result, and because all addresses are byte addresses with each successive byte having an address one greater than the preceding byte, the two LSBs of the class structure address are zero. These zero LSBs are not stored in header 220H. Therefore, the header has two bits not used for the address storage. These bits (header LSBs 0 and 1) are used for object locking. Bit 0, also called the L bit or the LOCK bit, is set to 1 when the object is locked. Bit 1, also called the W or WANT bit, is set to 1 when a thread is blocked waiting to acquire the lock for object 220.

Tables A and B set forth exemplary pseudocode for circuitry that executes lock and unlock instructions for one embodiment of processor 110. That circuitry is part of execution unit 136 and/or execution control and microcode unit 166. The pseudocode language of Tables A and B is similar to the hardware description language Verilog® described, for example, in D.E. Thomas, J.P. Moorby, "The Verilog® Hardware Description Language" (1991). The pseudocode can be easily converted to Verilog, and the corresponding circuitry can be implemented using methods known in the art.

Table A shows pseudocode for a lock instruction. At each of steps 1-0 through 1-3 in Table A, the contents of the corresponding register LOCKADDR0 through LOCKADDR3 are compared with the address of the object to be locked. If there is a match, the corresponding register LOCKCOUNT is incremented (steps 1-0a, 1-1a, 1-2a, 1-3a) and compared with zero (steps 1-0b, 1-1b, 1-2b, 1-3b). If the LOCKCOUNT register becomes 0 after incrementation, an overflow has occurred, and a trap LockCountOverflowIncrementTrap is generated. Generation of a trap terminates execution of the instruction. If the trap is enabled, processor 110 starts executing a trap handler defined for the trap. As is well known in the art, a trap handler is a predefined set of computer code which executes when the trap occurs.

In some embodiments, the trap handler for LockCountOverflowIncrementTrap maintains a wider lock counter mLOCKCOUNT (FIG. 3) than the LOCKCOUNT register. More particularly, in some embodiments, the operating system keeps track of locked objects using tables 310 in memory 120. A separate table 310 is kept for each thread. A table 310 is created for a thread when the thread is created, and the table 310 is deallocated when the corresponding thread is destroyed. Each table 310 includes a number of entries (mLOCKADDR, mLOCKCOUNT). The function of each entry is similar to the function of a register pair LOCKADDR/

- 5 -

LOCKCOUNT. More particularly, mLOCKADDR holds the address of an object locked by the thread. mLOCKCOUNT holds the count of lock instructions issued by the thread for the object. The count of lock instructions is the number of the lock instructions for which a corresponding unlock instruction has not been executed. If some mLOCKADDR = 0, then the entry is unused.

5           A table 310 may have more than four entries. Different tables 310 may have different numbers of entries.

Each memory location mLOCKADDR is 32 or 30 bits wide in some embodiments. Each location mLOCKCOUNT is 8 or more bits wide. In some embodiments, each location mLOCKCOUNT is 32 bits wide, and each register LOCKCOUNT is 8 bits wide.

10           When the operating system schedules a thread for execution, the operating system may load up to four entries from the corresponding table 310 into register pairs LOCKADDR/LOCKCOUNT. Each entry is written into a single register pair LOCKADDR/LOCKCOUNT. If mLOCKCOUNT is wider than LOCKCOUNT, the operating system writes to LOCKCOUNT as many LSBs of mLOCKCOUNT as will fit into LOCKCOUNT (8 LSBs in some embodiments). If some register pair does not receive an entry from table 310, the operating  
15           system sets the corresponding register LOCKADDR to 0 to indicate that the register pair is unused ("empty").

In some embodiments, table 310 includes a bit (not shown) for each entry to indicate whether the entry is to be written into a LOCKADDR/LOCKCOUNT register pair when the thread is scheduled for execution. In other embodiments, for each thread the operating systems keeps a list (not shown) of entries to be written to registers 144 when the thread is scheduled for execution. In some embodiments, the operating system has a bit  
20           for each entry, or a list of entries, to mark entries that have been written to LOCKADDR/LOCKCOUNT registers.

In some cases, lock and unlock instructions do not cause a trap to be generated. Therefore, the mLOCKCOUNT LSBs may be invalid, or there may be no entry in a table 310 for a lock specified by a LOCKADDR/LOCKCOUNT register pair.

25           When some thread T1 is preempted and another thread T2 is scheduled for execution on processor 110, the operating system writes all the non-empty LOCKADDR/LOCKCOUNT register pairs to the table 310 of thread T1 before loading the registers from the table 310 of thread T2. If mLOCKCOUNT is wider than LOCKCOUNT, the operating system writes each LOCKCOUNT register to the LSBs of the corresponding location mLOCKCOUNT. If the current thread's table 310 does not have an entry for a lock specified by a  
30           LOCKADDR/LOCKCOUNT register pair, an entry is created by the operating system.

In some embodiments, the trap handler for LockCountOverflowTrap searches the table 310 of the current thread for the entry with mLOCKADDR containing the address of the object to be locked. If such an

- 6 -

entry does not exist, the trap handler finds a free entry, and sets its mLOCKADDR to the address of the object to be locked and mLOCKCOUNT to zero. In either case (whether the entry existed or has just been created), the trap handler increments the mLOCKCOUNT MSBs which are not stored in the LOCKCOUNT register, and sets the LSBs to zero.

5 We now return to describing execution of the lock instruction by execution unit 136. In some embodiments, the comparisons of the registers LOCKADDR with the address of the object to be locked at steps 1-0 through 1-3 of Table A are performed in parallel by four comparators corresponding to the four registers, and the incrementation of LOCKCOUNT at steps 1-0a, 1-1a, 1-2a, 1-3a is performed using incrementors. Such comparators and incrementors are known in the art.

10 Execution unit 136 reads the LOCK bit (FIG. 2) from the header 220H of the object to be locked, and sets the LOCK bit to 1 to indicate that the object is locked (step 2a). This read-and-set (test-and-set) operation is an atomic operation, that is, (1) the processor will not take an interrupt until the operation is completed, and (2) in a multiprocessor environment, no other processor will be able to access the LOCK bit until the operation is completed. In some embodiments, this test-and-set operation is done in parallel with steps 1-0 through 1-3. In  
15 other embodiments, this test-and-set operation is done after steps 1-0 through 1-3, and only if none of the LOCKADDR registers contains the address of the object to be locked.

If none of the LOCKADDR registers contains the address of the object to be locked (step 2), and the LOCK bit was set before the test-and-set operation (step 2a), processor 110 generates a trap LockBusyTrap.

The trap handler for LockBusyTrap searches the table 310 of the current thread to see if the current  
20 thread holds the lock for the object. If the object address equals an address stored in mLOCKADDR in one of the entries of the table 310, the corresponding mLOCKCOUNT is incremented by the trap handler. Additionally, in some embodiments the trap handler may place the entry into a register pair LOCKADDR/LOCKCOUNT. This is desirable if the next lock or unlock instruction to be issued by the thread is likely to be for the object for which the thread issued the most recent lock instruction. If the trap handler desires to place the entry into a  
25 register pair but all the register pairs are taken by other locks, the trap handler vacates one of the register pairs by writing the register pair to the table 310. (The LOCKCOUNT register is written to the mLOCKCOUNT LSBs if mLOCKCOUNT is wider than LOCKCOUNT, as described above.)

If the current thread does not hold the lock and thus the object address does not match any of the memory locations mLOCKADDR in the corresponding table 310, the trap handler sets the WANT bit in the  
30 object header (FIG. 2) and places the thread into a queue of threads waiting to acquire this lock.

We return now to describing the execution of the lock instruction by execution unit 136. If the object's LOCK bit was not set before the test-and-set operation, steps 2b-0 through 2b-3 are executed. At each step 2b-i (i = 0 through 3), a respective comparator compares the register LOCKADDR<sub>i</sub> with zero. This comparison is

- 7 -

performed in parallel with comparisons of steps 1-0 through 1-3 and 2. If LOCKADDR0 = 0 (step 2b-0), the register pair LOCKADDR0/LOCKCOUNT0 is unused. Register LOCKADDR0 is written with the address of the object being locked (step 2b-0a). The register LOCKCOUNT0 is set to 1 (step 2b-0b).

5 If LOCKADDR0 is not 0 but LOCKADDR1 = 0, then register LOCKADDR1 is written with the address of the object to be locked, and register LOCKCOUNT1 is set to 1 (steps 2b-1a, 2b-1b). If LOCKADDR0 and LOCKADDR1 are not 0 but LOCKADDR2 = 0, then LOCKADDR2 is written with the address of the object to be locked, and register LOCKCOUNT2 is set to 1 (steps 2b-2a, 2b-2b). If LOCKADDR0, LOCKADDR1, and LOCKADDR2 are not 0 but LOCKADDR3 = 0, then register LOCKADDR3 is written with the address of the object to be locked, and register LOCKCOUNT3 is set to 1  
10 (steps 2b-3a, 2b-3b).

If none of the LOCKADDR registers is equal to 0, then the trap NoLockAddrRegsTrap is generated (step 2c). In some embodiments, the trap handler for this trap finds or creates a free entry in the table 310 of the current thread. The trap handler writes the address of the object to be locked into location mLOCKADDR of that entry, and sets the corresponding mLOCKCOUNT to 1. Additionally, the trap handler may place the table entry  
15 into a LOCKADDR/LOCKCOUNT register pair. The old contents of the register pair are stored in the thread's table 310 before the register pair is written. It will be appreciated that other types of known replacement algorithms may also be used, e.g., a least recently used type of replacement algorithm or a random type of replacement algorithm.

Table B shows exemplary pseudocode for the unlock instruction. At steps 1-0 through 1-3, the  
20 LOCKADDR registers are compared in parallel with the address of the object to be unlocked. If a match occurs, this indicates that the current thread holds the lock, and the corresponding LOCKCOUNT register is decremented by a decrementor (steps 1-0a, 1-1a, 1-2a, 1-3a) and compared with zero (steps 1-0b, 1-1b, 1-2b, 1-3b). If the LOCKCOUNT register becomes 0 after decrementation, the trap LockCountZeroDecrementTrap is generated. As described above, in some embodiments, the locations mLOCKCOUNT in tables 310 are wider than the  
25 LOCKCOUNT register. In some such embodiments, the trap handler for LockCountZeroDecrementTrap searches the corresponding table 310 for an entry whose mLOCKADDR stores the address of the object being unlocked. If such entry is found, the trap handler checks the mLOCKCOUNT location corresponding to the LOCKCOUNT register which was decremented to 0. If that mLOCKCOUNT location has a "1" in the MSBs that were not written into the LOCKCOUNT register, the object remains locked by the thread. In the  
30 mLOCKCOUNT memory location the field formed by the MSBs is decremented, and the LSBs are set to 11...1 (all 1's) and are written to the LOCKCOUNT register.

If the mLOCKCOUNT MSBs are all 0's, or if there is no entry with mLOCKADDR holding the address of the object being unlocked, then the trap handler frees the lock making it available for other threads. Freeing the lock is described in more detail below.



- 8 -

If the mLOCKCOUNT locations are not wider than the LOCKCOUNT registers, the trap handler need not check an mLOCKCOUNT location to determine whether the lock is to be freed.

Freeing the lock involves the following operations. The trap handler examines the WANT bit of object header 220H. If the WANT bit is set, another thread is blocking on this lock. The trap handler selects one of such threads, sets its status to runnable, and gives the lock to this thread. In particular, the trap handler writes the count of 1 into the LOCKCOUNT register. If there was a corresponding pair mLOCKADDR/mLOCKCOUNT, the trap handler writes 1 to the mLOCKCOUNT location. Alternatively, in some embodiments, the trap handler writes 0 to the mLOCKADDR location to deallocate the mLOCKADDR/mLOCKCOUNT pair. Further, if the thread receiving the lock is the only thread that has been blocking on the lock, the trap handler resets the WANT bit.

If there were no threads blocking on the lock, the trap handler writes zero to (a) the corresponding LOCKADDR register and (b) the corresponding mLOCKADDR location if one exists. In addition, the trap handler resets the LOCK bit in header 220H. Also, if the current thread's table 310 includes a non-empty entry which could not be written into the LOCKADDR/LOCKCOUNT registers because the registers were unavailable, the trap handler places one of the entries into the LOCKADDR/LOCKCOUNT register pair which is being vacated by the lock freeing operation.

If none of the LOCKADDR registers holds the address of the object to be unlocked (step 2), the LockReleaseTrap is generated. The associated trap handler searches the mLOCKADDR locations of the current thread's table 310 for the address of the object to be unlocked. If a match occurs, the corresponding location mLOCKCOUNT is decremented by the trap handler. If mLOCKCOUNT becomes 0, the lock is freed. To free the lock, the trap handler perform operations similar to those described above for the trap LockCountZeroDecrementTrap. More particularly, if the WANT bit is set, the trap handler finds another thread blocking on the lock and sets that thread's status to runnable. The trap handler sets the corresponding location mLOCKCOUNT to 1. In some embodiments, the trap handler places the mLOCKADDR/mLOCKCOUNT entry into a LOCKADDR/LOCKCOUNT register pair. If the thread receiving the lock is the only thread that has been blocking on the lock, the trap handler resets the WANT bit. If there were no threads blocking on the lock (the WANT bit was 0), the trap handler writes zero to the mLOCKADDR location and resets the LOCK bit in object header 220H.

If none of the memory locations mLOCKADDR in table 310 of the current thread holds the address of the object to be unlocked, the trap handler generates the exception IllegalMonitorStateException. In some embodiments, this exception is a Java™ throw. More particularly, in some embodiments, processor 110 executes Java™ Virtual Machine language instructions (also known as Java byte codes). The Java™ Virtual Machine

language is described, for example, in T. Lindholm and F. Yellin, "The Java™ Virtual Machine Specification" (1997) incorporated herein by reference.

Processor 110 provides fast locking and unlocking in many of the following common situations: when there is no contention for a lock, and when a thread performs multiple lock operations on the same object before the object lock is freed. More particularly, when a lock instruction is issued, in many cases the object has not  
5 been locked by another thread (that is, no contention occurs). If the object has already been locked by the same thread that has now issued the lock instruction, in many cases the address of the object is already in a LOCKADDR register because in many cases the thread does not hold more than four locks at the same time and all the locked object addresses for the thread are in the LOCKADDR registers. Even if not all the locked object  
10 addresses are in the LOCKADDR registers, there is a possibility that the address of the object specified by the lock instruction is in a LOCKADDR register. In many such cases, the locking operation requires incrementing the corresponding LOCKCOUNT register (Table A, steps 1-ia where  $i = 0, 1, 2, 3$ ), which is a fast operation in many embodiments. If the incrementation does not lead to an overflow, no trap will be generated.

Locking is also fast when the object has not been locked by any thread (including the thread issuing the  
15 lock instruction) if one of the register pairs LOCKADDR/LOCKCOUNT is unused. In such cases, the object is locked in one of steps 2b-0 through 2b-3 (Table A). Again, no trap is generated.

Similarly, in an unlock instruction, in many cases the address of the object to be unlocked will be in one of the LOCKADDR registers. If the corresponding LOCKCOUNT register is decremented to a non-zero value, no trap is generated.

20 In some embodiments, processor 110 is a microprocessor of type "picoJava I" whose specification is produced by Sun Microsystems of Mountain View, California. This microprocessor executes Java Virtual Machine instructions. The lock instruction is the "monitorenter" instruction of the Java Virtual Machine instruction set or the "enter\_sync\_method" instruction of the processor "picoJava I". The "enter\_sync\_method" instruction is similar to "monitorexit" but the "enter\_sync\_method" instruction takes as a parameter a reference to  
25 a method rather than an object. "Enter\_sync\_method" locks the receiving object for the method and invokes the method. The unlock instruction is the "monitorexit" instruction of the Java Virtual Machine instruction set or the return instruction from a method referenced in a preceding "enter\_sync\_method" instruction.

Some embodiments of processor 110 include more or less than four LOCKADDR/LOCKCOUNT register pairs.

30 In some embodiments, registers 144 include register triples (THREAD\_ID, LOCKADDR, LOCKCOUNT) as shown in FIG. 4. In each triple, the register THREAD\_ID identifies the thread which holds the lock recorded in the register pair LOCKADDR/LOCKCOUNT. When a lock or unlock instruction is issued,

- 10 -

execution unit 136 examines only those LOCKADDR/LOCKCOUNT pairs for which the register THREAD\_ID holds the ID of the current thread. In other respects, the execution of lock and unlock instructions is similar to the case of FIG. 2. The structure of FIG. 4 makes it easier to keep the locked objects' addresses and lock counts in registers 144 for different threads at the same time. In some embodiments used with the structure of FIG. 4, the operating system does not reload the registers 144 when a different thread becomes scheduled for execution. The operating system maintains a table 310 for each thread as shown in FIG. 3. When a register triple needs to be vacated, the corresponding LOCKADDR/LOCKCOUNT values are written to the corresponding table 310. When a table entry is placed into a register pair LOCKADDR/LOCKCOUNT, the corresponding register THREAD\_ID is written with the ID of the corresponding thread.

10           The processors of FIGs. 1-4 are suitable for efficient implementation of the Java Virtual Machine lock and unlock instructions "monitorenter" and "monitorexit". The counters associated with the object monitors in Java can be implemented using registers LOCKCOUNT.

In some embodiments, registers LOCKCOUNT and locations mLOCKCOUNT are omitted. The processor does not keep track of the lock counts, and the processor frees a lock on any unlock instruction corresponding to the lock. The processor operation is similar to the operation described above in connection with Tables A and B. However, in Table A, steps 1-0 through 1-3b are omitted. Steps 2b-0b, 2b-1b, 2b-2b, and 2b-3b (LOCKCOUNT operations) are also omitted. In Table B, step 1-0a is omitted, and at step 1-0b the trap LockCountZeroDecrementTrap is generated unconditionally. The same applies to steps 1-1a and 1-1b, 1-2a and 1-2b, 1-3a and 1-3b.

20           In some embodiments, each LOCKCOUNT register is 1-bit wide, and the processor frees a lock on any unlock instruction corresponding to the lock.

The monitorenter bytecode instruction of the Java™ VM requires LOCKCOUNTs. Thus, for a microprocessor that implements a Java™ Virtual Machine (i.e., a bytecode engine), it is advantageous to implement the monitorenter bytecode instruction of the Java™ VM by providing a lock cache that holds LOCKCOUNTs.

However, in an instruction set architecture of a general microprocessor, locking of computer resources may not require LOCKCOUNTs. Moreover, elimination of atomics and traps in thread synchronization is desired. Accordingly, an efficient and fast locking mechanism for a general microprocessor that checks whether a current thread executing on a processor owns a computer resource is provided. In one embodiment, an instruction for a microprocessor is provided to determine whether a current thread executing on a processor of the microprocessor owns a computer resource by checking a lock cache of the processor.

- 11 -

FIG. 5 is a block diagram illustrating lock caches 512 and 514 in accordance with one embodiment of the present invention. P1 processor 502 of multiprocessor 500 includes a lock cache 512 and an execution unit 516 executing, for example, thread\_A. P2 processor 504 includes a lock cache 514 and an execution unit 518 executing, for example, thread\_B. Memory 120 stores an object header 220H as similarly discussed above. In one embodiment, lock caches 512 and 514 each include four registers to provide four-entry caches, which are disposed on execution units 516 and 518, respectively.

In this embodiment, LOCKCOUNTs are not stored (e.g., cached) in the lock caches. The operation of multiprocessor 500 using lock caches 512 and 514 is discussed below with respect to FIGs. 7 and 8. The operation of maintaining memory 120 and loading the lock caches with the appropriate information from memory 120 upon beginning execution of a new thread can be similar to the implementation and use of memory 120 discussed above.

In particular, lock caches 512 and 514 cache an object address only if the current thread executing on the respective processor owns the object (e.g., has locked the object). No lock count is maintained in lock caches 512 and 514.

Referring to FIG. 6, a lockcachecheck instruction 602 includes an 8-bit lockcachecheck opcode 604, a 7-bit object\_address 606 (e.g., specified by a register specifier rs1), and a 7-bit register destination (Rd) 608 for the value returned by execution of the lockcachecheck instruction. For example, the lockcachecheck instruction can be executed on P1 processor 602 of FIG. 6 in one cycle which is performed in the E stage of the pipeline.

The execution of the lockcachecheck instruction performs the following operation: the object\_address is compared with each entry of lock cache 612 to determine if the object\_address, which corresponds to a computer resource, is held in the lock cache (i.e., a cache hit/miss is determined). In a preferred embodiment, lock cache is implemented with a content addressable memory (CAM) or other associative memory because the memory determines cache hits in parallel. The result of the hit/miss operation is returned to Rd 608. For example, Rd is set to 0 if the lockcachecheck operation results in a miss, and Rd is set to 1 if the lockcachecheck operation results in a hit. One of ordinary skill in the art will recognize that there are various way to implement the circuitry for performing the operation of the lockcachecheck instruction in a microprocessor, such as the microprocessor of FIG. 5.

Referring to FIG. 7, a lockcachecheck instruction is executed on P1 processor 602 of FIG. 6 at stage 702. At stage 704, the lockcachecheck instruction returns 0 (i.e., Rd = 0) if the object (i.e., address\_of\_object) is not cached in lock cache 712 (i.e., a cache miss). The lockcachecheck instruction returns 1 (i.e., Rd = 1) if the object (i.e., address\_of\_object) is cached in lock cache 712 (i.e., a cache hit). The operation of stage 704 can be implemented in hardware in a microprocessor, such as P1 processor 602 of FIG. 6. At stage 706, if Rd = 0, then a branch operation, for example, is taken to attempt to acquire the lock. Thus, software is responsible for

- 12 -

handling a lock cache miss (e.g., the software can determine if the current thread owns the object by checking a table maintained in memory, and if not, then the software can branch to a loop that waits for the object to become available). Otherwise (i.e., Rd = 1), P1 processor 602 continues executing the current instruction stream, because the current thread executing on processor P1 602 owns the object. If the lock count is not incremented in hardware, then software may increment the lock count.

Accordingly, in the embodiments described above with respect to FIGs. 5 through 7, elimination of atomics and traps in thread synchronization is provided. For example, a lockcachecheck instruction can be provided for a general microprocessor in which it is desirable to provide for an efficient implementation of locking of computer resources.

The above embodiments illustrate but do not limit the present invention. The present invention is not limited by any particular processor architecture, the presence or structure of caches or memory, or the number of bits in any register or memory location. The present invention is not limited to any particular types of objects that can be locked or unlocked. An object can represent any computer resource, including such resources as data, critical code sections, hardware, or any combination of the above. Some embodiments create an object dedicated to represent a computer resource for locking and unlocking operations. While in embodiments described above an unused register LOCKADDR is identified by zero in the LOCKADDR register, in some embodiments an unused register is identified by some non-zero value in the LOCKADDR register, by some value in the THREAD\_ID register, by a separate bit, or some combination thereof. A similar statement is true for unused mLOCKADDR locations. In some embodiments, some operations described above as performed by hardware are performed by software instead of hardware. The present invention is not limited to addresses being byte addresses. Other embodiments and variations are within the scope of the present invention, as defined by the appended claims.

TABLE A  
Lock Instruction

```

1-0.  if (LOCKADDR0 == address of object to be locked)
      {
1-0a.  LOCKCOUNT0++;
1-0b.  if (LOCKCOUNT0 == 0) /* LOCKCOUNT0 overflowed */
      LockCountOverflowIncrementTrap;
      }

1-1.  if (LOCKADDR1 == address of object to be locked)
      {
1-1a.  LOCKCOUNT1++;
1-1b.  if (LOCKCOUNT1 == 0) /* LOCKCOUNT1 overflowed */
      LockCountOverflowIncrementTrap;
      }

1-2.  if (LOCKADDR2 == address of object to be locked)
      {
1-2a.  LOCKCOUNT2++;
1-2b.  if (LOCKCOUNT2 == 0) /* LOCKCOUNT2 overflowed */
      LockCountOverflowIncrementTrap;
      }

1-3.  if (LOCKADDR3 == address of object to be locked)
      {
1-3a.  LOCKCOUNT3++;
1-3b.  if (LOCKCOUNT3 == 0) /* LOCKCOUNT3 overflowed */
      LockCountOverflowIncrementTrap;
      }

2.    if (none of LOCKADDR0, LOCKADDR1, LOCKADDR2, LOCKADDR3 is equal to address of object
      to be locked)
      {
2a.    Test the LOCK bit in the object header, and set the LOCK bit to 1. (This test-and-set
      operation is an atomic operation.)
      if (the LOCK bit was set before the test-and-set operation)
      LockBusyTrap;

2b-0.  else if (LOCKADDR0 == 0) /* LOCKADDR0 unused */
      {
2b-0a.  LOCKADDR0 = address of object to be locked;
2b-0b.  LOCKCOUNT0 = 1;
      }

2b-1.  else if (LOCKADDR1 == 0) /* LOCKADDR1 unused */
      {
2b-1a.  LOCKADDR1 = address of object to be locked;
2b-1b.  LOCKCOUNT1 = 1;
      }

2b-2.  else if (LOCKADDR2 == 0) /* LOCKADDR2 unused */
      {
2b-2a.  LOCKADDR2 = address of object to be locked;

```

- 14 -

```
2b-2b.          LOCKCOUNT2 = 1;
                  }

2b-3.           else if (LOCKADDR3 == 0) /* LOCKADDR3 unused */
                  {
2b-3a.          LOCKADDR3 = address of object to be locked;
2b-3b.          LOCKCOUNT3 = 1;
                  }
2c.             else NoLockAddrRegsTrap;
                  }
```

- 15 -

**TABLE B**  
Unlock Instruction

- 1-0. if (LOCKADDR0 == address of object to be unlocked)  
     {  
     1-0a. LOCKCOUNT0--;  
     1-0b. if (LOCKCOUNT0 == 0)  
             LockCountZeroDecrementTrap;  
     }
- 1-1. if (LOCKADDR1 == address of object to be unlocked)  
     {  
     1-1a. LOCKCOUNT1--;  
     1-1b. if (LOCKCOUNT1 == 0)  
             LockCountZeroDecrementTrap;  
     }
- 1-2. if (LOCKADDR2 == address of object to be unlocked)  
     {  
     1-2a. LOCKCOUNT2--;  
     1-2b. if (LOCKCOUNT2 == 0)  
             LockCountZeroDecrementTrap;  
     }
- 1-3. if (LOCKADDR3 == address of object to be unlocked)  
     {  
     1-3a. LOCKCOUNT3--;  
     1-3b. if (LOCKCOUNT3 == 0)  
             LockCountZeroDecrementTrap;  
     }
2. if (none of LOCKADDR0, LOCKADDR1, LOCKADDR2, LOCKADDR3 is equal to address of object to be unlocked)  
     LockReleaseTrap



**WE CLAIM**

- 1           1.       An apparatus, comprising:  
2           a lock cache, the lock cache holding values that identify locks for computer resources; and  
3           logic circuitry, the logic circuitry connected to the lock cache, and the logic circuitry performing an  
4           operation that determines whether a computer resource is held in the lock cache by receiving a  
5           value V1 identifying the computer resource and then determining whether the lock cache holds  
6           the value V1.
- 1           2.       The apparatus of Claim 1 wherein the logic circuitry comprises:  
2           a comparator for comparing the value V1 with a value V2 held in the lock cache.
- 1           3.       The apparatus of Claim 1 wherein the logic circuitry returns a first predetermined value if the  
2           lock cache is not holding the value V1.
- 1           4.       The apparatus of Claim 1 wherein the logic circuitry returns a second predetermined value if  
2           the lock cache is holding the value V1.
- 1           5.       The apparatus of Claim 1 wherein the logic circuitry comprises a computer processor, and the  
2           lock cache comprises registers of an execution unit of the computer processor.
- 1           6.       A process, comprising:  
2           receiving a value V1 identifying a computer resource; and  
3           determining if a lock cache of a computer processor holds the value V1, wherein the lock cache holding  
4           the value V1 indicates that the computer processor owns the computer resource.
- 1           7.       The process of Claim 6 further comprising:  
2           returning a first predetermined value if no entry of the lock cache holds the value V1.
- 1           8.       The process of Claim 7 further comprising:  
2           returning a second predetermined value if the lock cache holds the value V1.
- 1           9.       The process of Claim 7 wherein the process comprises execution of an instruction on the  
2           computer processor of a microprocessor.

- 17 -

- 1           10.     An apparatus, comprising:  
2           a lock cache of a computer processor of a microprocessor, the lock cache holding values that identify  
3           computer resources owned by the computer processor; and  
4           means for performing a lock cache check operation that determines whether a value V1 corresponding  
5           to a computer resource is held in the lock cache.
- 1           11.     The apparatus of Claim 10 wherein the means for performing a lock cache check operation  
2     further comprises:  
3           means for comparing the value V1 with each value held in the lock cache.
- 1           12.     The apparatus of Claim 11 wherein the means for performing a lock cache check operation  
2     further comprises:  
3           means for returning a first predetermined value if the lock cache is not holding the value V1.
- 1           13.     The apparatus of Claim 12 wherein the means for performing a lock cache check operation  
2     further comprises:  
3           means for returning a second predetermined value if the lock cache is holding the value V1.
- 1           14.     The apparatus of Claim 13 wherein the means for performing a lock cache check operation  
2     comprises the microprocessor executing a lock cache check instruction.
- 1           15.     The apparatus of Claim 14 wherein the lock cache comprises four entries.
- 1           16.     The apparatus of Claim 14 wherein the lock cache comprises registers of an execution unit of  
2     the processor.
- 1           17.     The apparatus of Claim 10 wherein the lock cache holds the value V1 that corresponds to the  
2     computer resource.
- 1           18.     The apparatus of Claim 17 wherein the computer resource comprises an object.
- 1           19.     The apparatus of Claim 18 wherein the object comprises a data structure.
- 1           20.     The apparatus of Claim 10 wherein the lock cache comprises object addresses, the object  
2     addresses being addresses of objects owned by a thread executing on the processor.

1/5

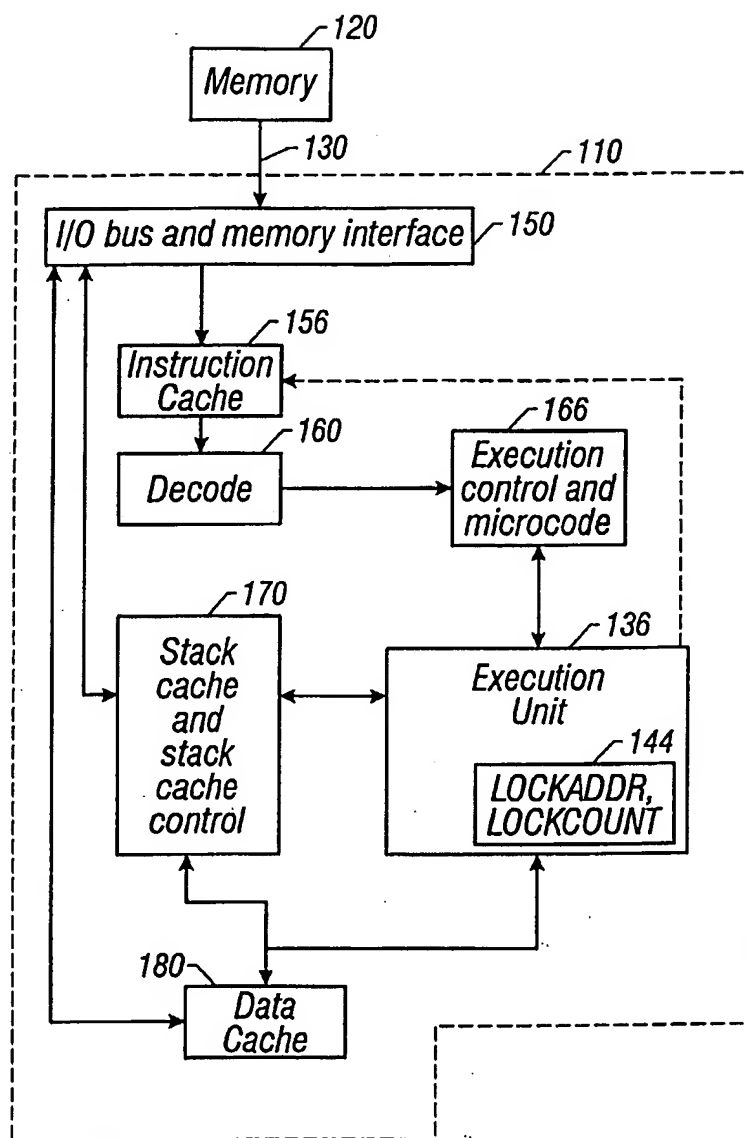


FIG. 1

2/5

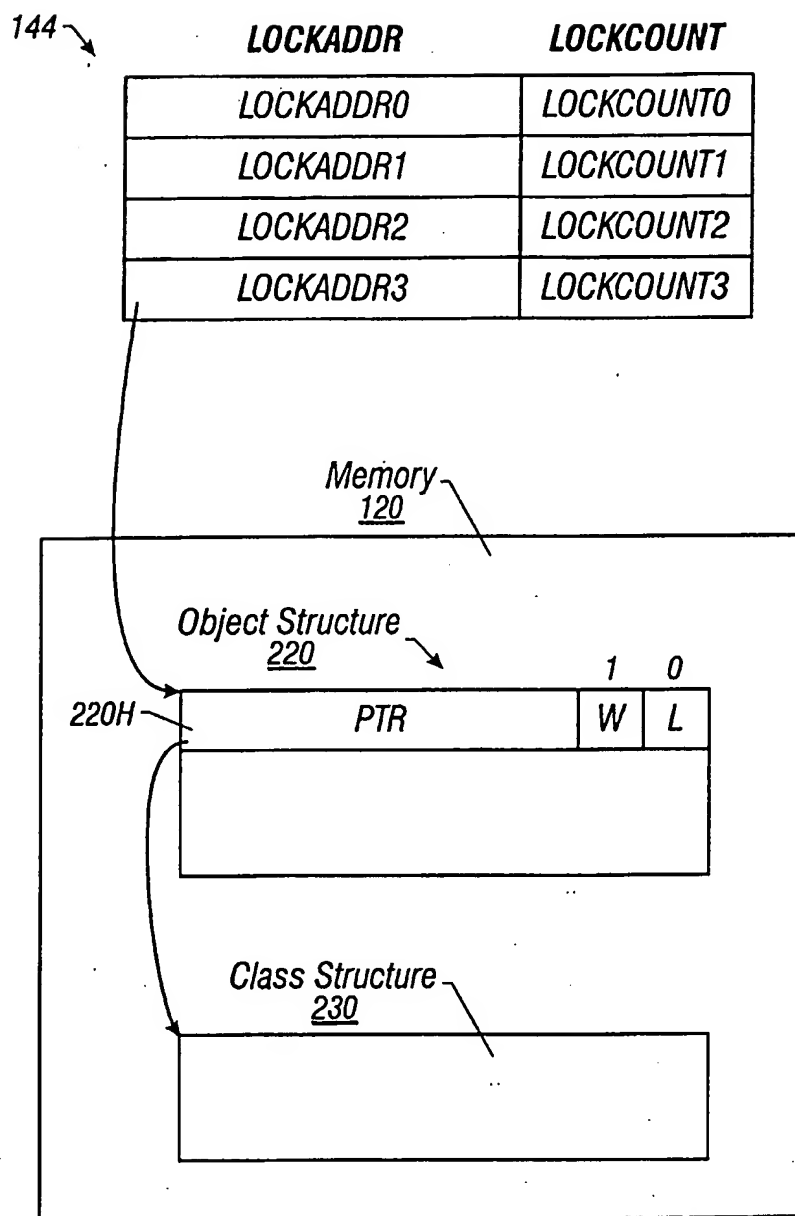


FIG. 2

3/5

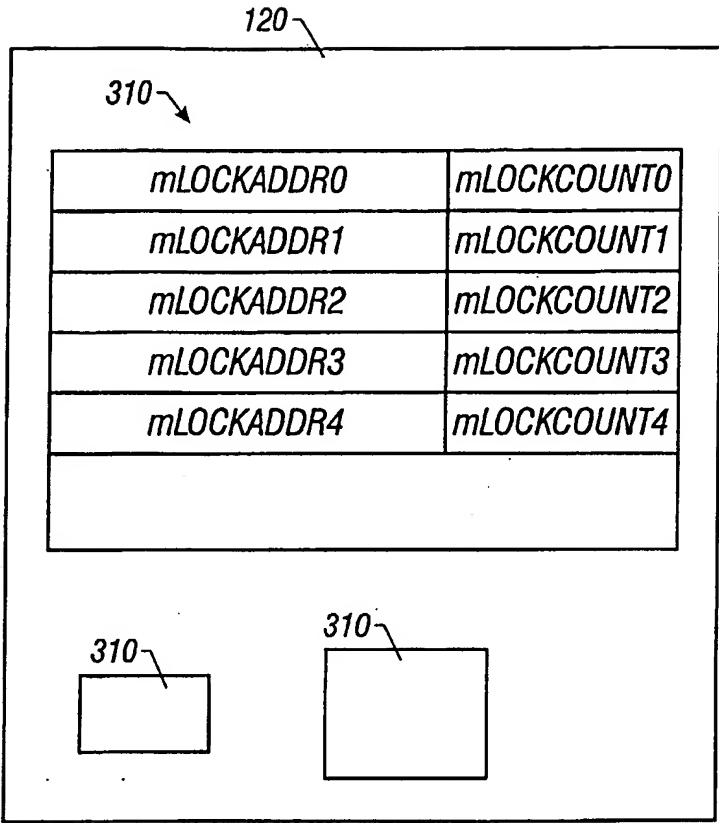


FIG. 3

144

THREAD_ID	LOCKADDR	LOCKCOUNT
THREAD_ID0	LOCKADDR0	LOCKCOUNT0
THREAD_ID1	LOCKADDR1	LOCKCOUNT1
THREAD_ID2	LOCKADDR2	LOCKCOUNT2
THREAD_ID3	LOCKADDR3	LOCKCOUNT3

FIG. 4

4/5

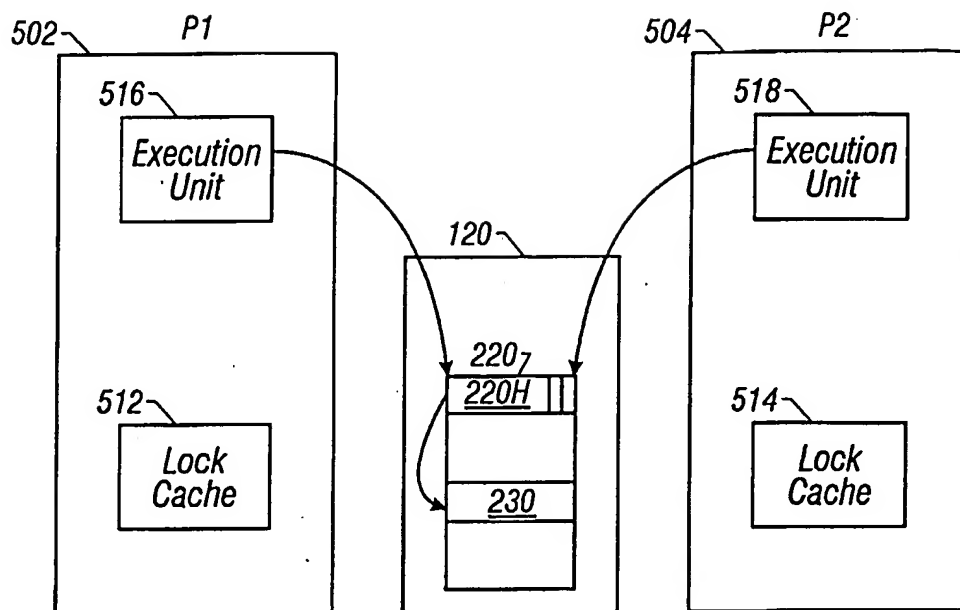


FIG. 5

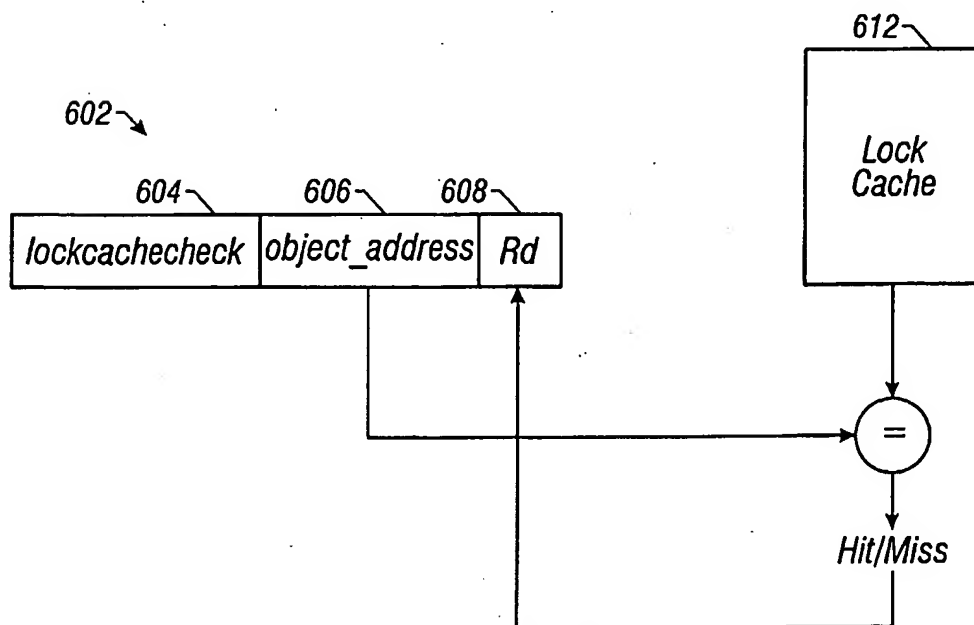


FIG. 6

5/5

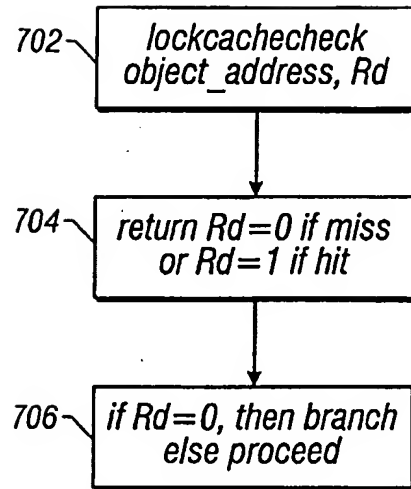


FIG. 7

## INTERNATIONAL SEARCH REPORT

In application No  
PCT/US 99/28875A. CLASSIFICATION OF SUBJECT MATTER  
IPC 7 G06F12/08 G06F9/46

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X Y A	US 5 404 482 A (STAMM REBECCA L ET AL) 4 April 1995 (1995-04-04)  abstract column 3, line 30 -column 4, line 22 column 10, line 65 -column 12, line 26 column 28, line 17 - line 43 ---	1-4, 6-15,17 5,16 18-20
X Y A	US 5 555 393 A (TANAKA TETSUYA ET AL) 10 September 1996 (1996-09-10)  abstract column 2, line 3 - line 46 column 3, line 55 -column 4, line 32 --- -/-	1-4,6-8, 10-13 5 9,14-20

☒ Further documents are listed in the continuation of box C.☒ Patent family members are listed in annex.

## \* Special categories of cited documents:

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

3 May 2000

Date of mailing of the international search report

11/05/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

Archontopoulos, E



## INTERNATIONAL SEARCH REPORT

International Application No  
PCT/US 99/28875

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	EP 0 567 355 A (DIGITAL EQUIPMENT CORP) 27 October 1993 (1993-10-27) page 8, line 26 -page 9, line 40 ---	5,16
A	US 5 615 167 A (JAIN ANIL K ET AL) 25 March 1997 (1997-03-25) abstract column 2, line 35 - line 45 column 4, line 24 -column 5, line 25 claims 1,3 ---	1,6,10
A	FR 2 733 067 A (NIPPON ELECTRIC CO) 18 October 1996 (1996-10-18) abstract page 4, line 25 -page 5, line 2 page 9, line 22 -page 11, line 18 -----	5,16

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 99/28875

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5404482	A	04-04-1995	US 5155843 A	13-10-1992
			CA 2045808 A	30-12-1991
			DE 69127936 D	20-11-1997
			DE 69127936 T	07-05-1998
			EP 0468831 A	29-01-1992
			JP 6095963 A	08-04-1994
			US 5317720 A	31-05-1994
			US 5432918 A	11-07-1995
			US 5404483 A	04-04-1995
			US 5347648 A	13-09-1994
			US 5471591 A	28-11-1995
			CA 2045715 A	30-12-1991
			DE 69111707 D	07-09-1995
			DE 69111707 T	28-03-1996
			EP 0465319 A	08-01-1992
			JP 6095964 A	08-04-1994
-----				
US 5555393	A	10-09-1996	US 5487162 A	23-01-1996
-----				
EP 0567355	A	27-10-1993	US 5553258 A	03-09-1996
			US 5553266 A	03-09-1996
-----				
US 5615167	A	25-03-1997	NONE	
-----				
FR 2733067	A	18-10-1996	JP 2776759 B	16-07-1998
			JP 8286968 A	01-11-1996
			US 5893156 A	06-04-1999
-----				